Packet Pair Layered Multicast Congestion Control

Arnaud Legout

Eurecom

Arnaud.Legout@eurecom.fr

Motivation

- Audio/video applications take a growing place in the Internet.
- No satisfactory CC protocols for multicast delivery:
  - RLM (low convergence time, unstability, unfairness, loss induced, ...).
  - TCP-friendly versions (low convergence time, loss induced, ...).
  - Both need fine tuning of parameters (unresolved issue).
- We apply the FS-paradigm to devise a new CC protocol for audio/video applications and multicast delivery.
- The Fair Scheduler (FS) paradigm (set of assumptions):
  - Network Part (NP): We assume a Fair Scheduler network.
  - End System Part (ESP): We assume selfish and non-collaborative end users.

Implications of the FS-paradigm

- FS-paradigm for the design of CC protocols:
  - No need for specific mechanisms to improve one of the properties of an ideal CC protocol.
  - Just address the application needs.
- The FS-paradigm does not give the mechanisms to meet the application needs but considerably simplifies the design of CC protocols.
- We devise PLM according to the FS-paradigm, we do not specifically address the properties of an ideal CC protocol.

PLM Principle

- Assumptions:
  - Data that can be stripped in cumulative layers (mainly audio/video).
  - Multicast capable network.
  - Fair scheduler network.
- PLM scheme:
  - Receiver-driven.
  - Cumulative layers.
  - **The source sends on each layer packet by pair (PP).**
- The PPs allow to dynamically infer the available bandwidth for each receiver.

PLM Algorithm

- Each PP received leads to an estimate of the available bandwidth.
- We drop layers each time we have an estimate lower than the current layer subscription until the layer subscription is lower than the estimate).
- We add layers according to the minimum estimate received during a period C if all the estimates received during C are greater than the current layer subscription.

Simulations: Basic Scenarios

- A single PLM session with a simple topology with a large heterogeneity of bandwidth and delay:
  - We evaluate the speed, the stability, and the accuracy of the convergence of PLM.
- A single PLM session with a large number of receivers and a single bottleneck:
  - We evaluate the scaling properties of PLM with a large number of receivers and with late join.
- 3 PLM sessions and 3 CBR flows through a single bottleneck:
  - We evaluate the adaptation of PLM with bottleneck change and the basic behavior of multiple PLM sessions.

A Single PLM Session: Convergence

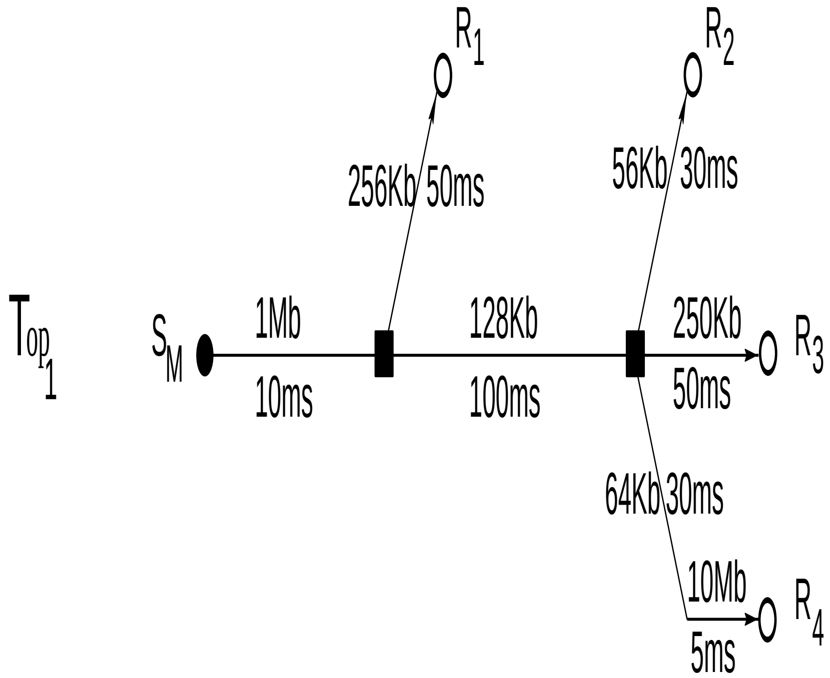

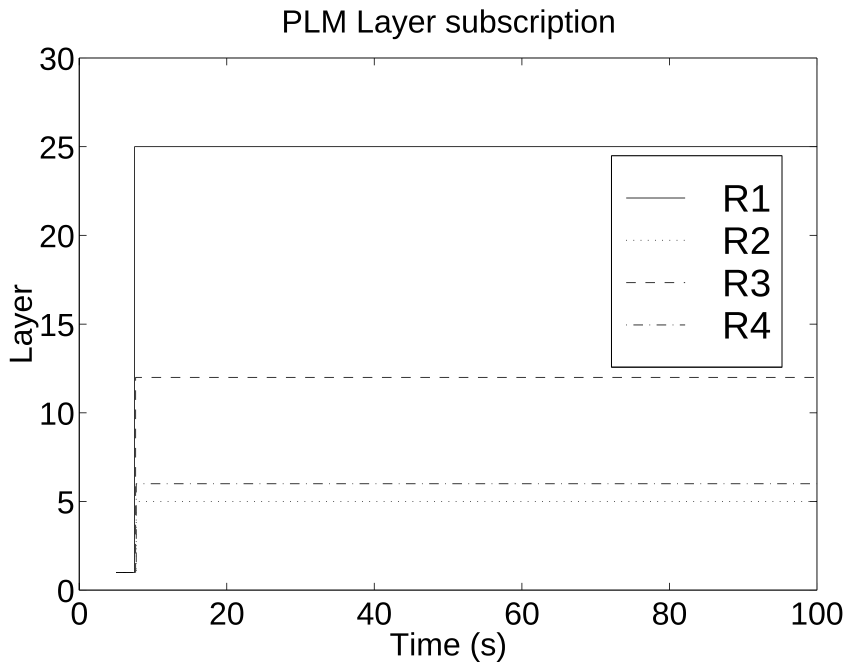

- Evaluation of the speed and the accuracy of the convergence in the context of a large heterogeneity of delay and bandwidth.
- 10Kbit/layer.

- All the receivers converge to the optimal rate in the order of C=1 second and stay at this rate during the whole simulation. **No loss induced**.

A Single RLM Session: Convergence

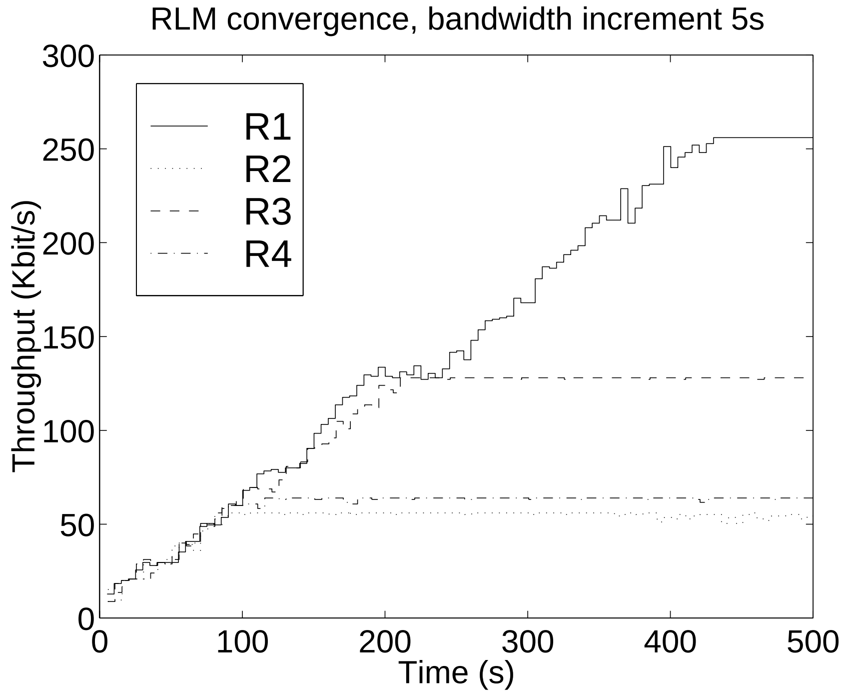

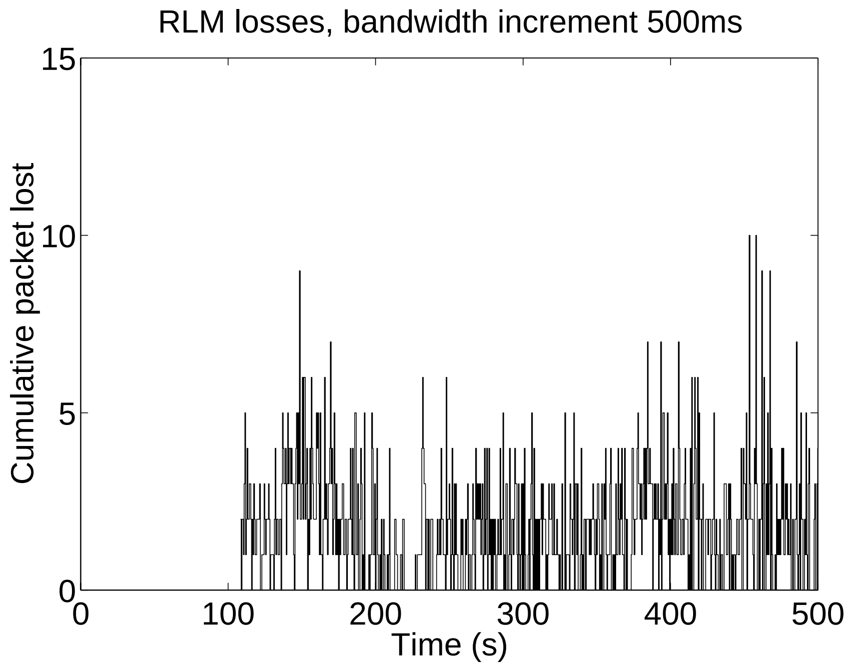

- Must run 500s. Low convergence time.
- Significant number of losses induced (3%), due to the loss threshold (25%).

A Single PLM Session: Scalability

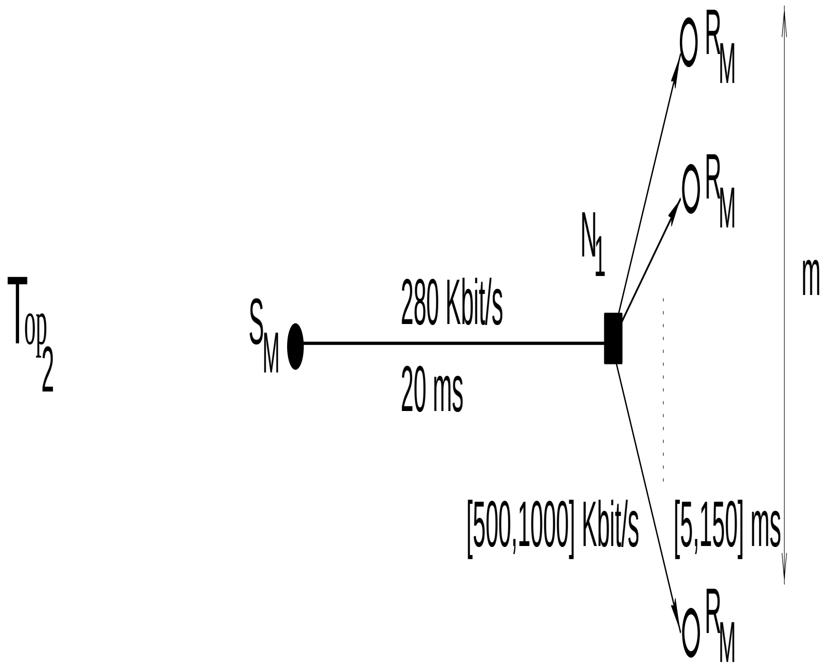

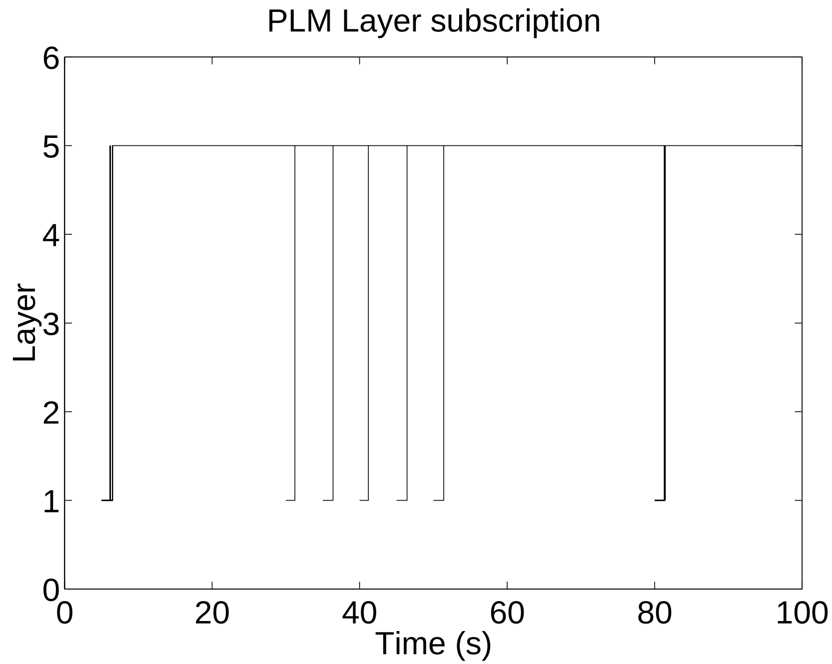

- Evaluation of the scalability properties of PLM with the number of receivers.
- 50Kbit/layer.
- 20+5+5 receivers.
- PLM convergence is independent of the number of receivers and of the late joins. **No loss induced**.

A Single RLM Session: Scalability

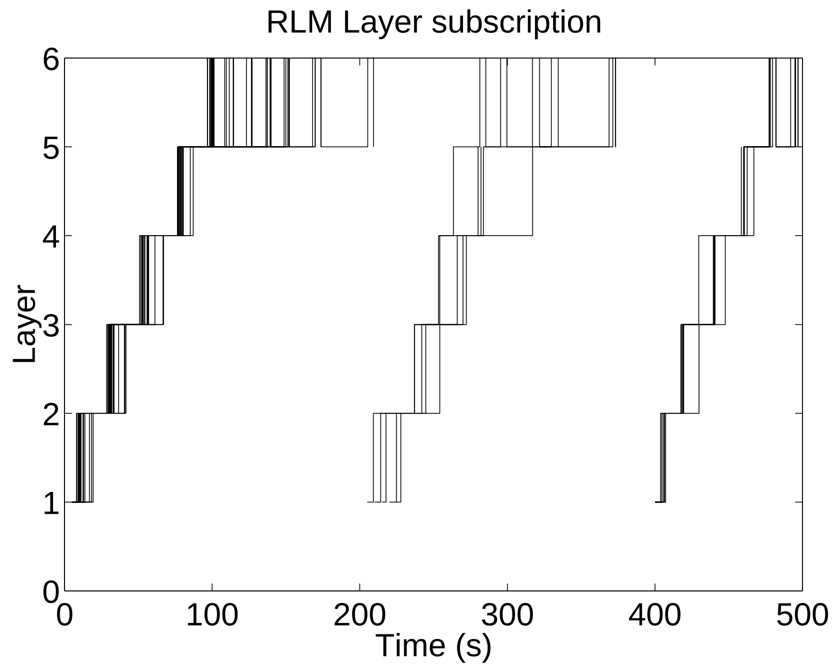

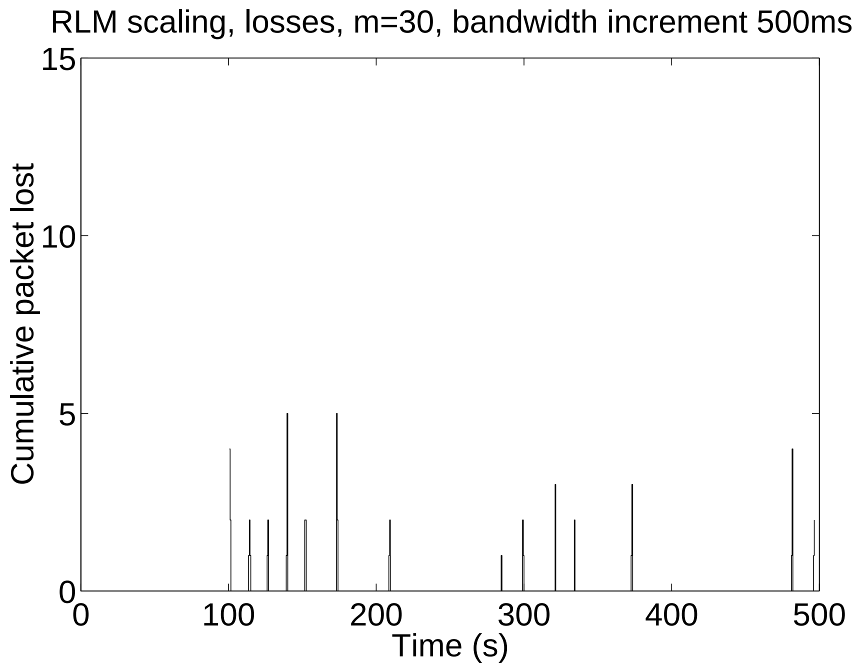

- Must run 500s. Low convergence time. Synchronization of the joins due to the shared learning.
- Some losses induced (only due to the join experiments, loss<0,008%).

Multiple PLM and CBR Sessions

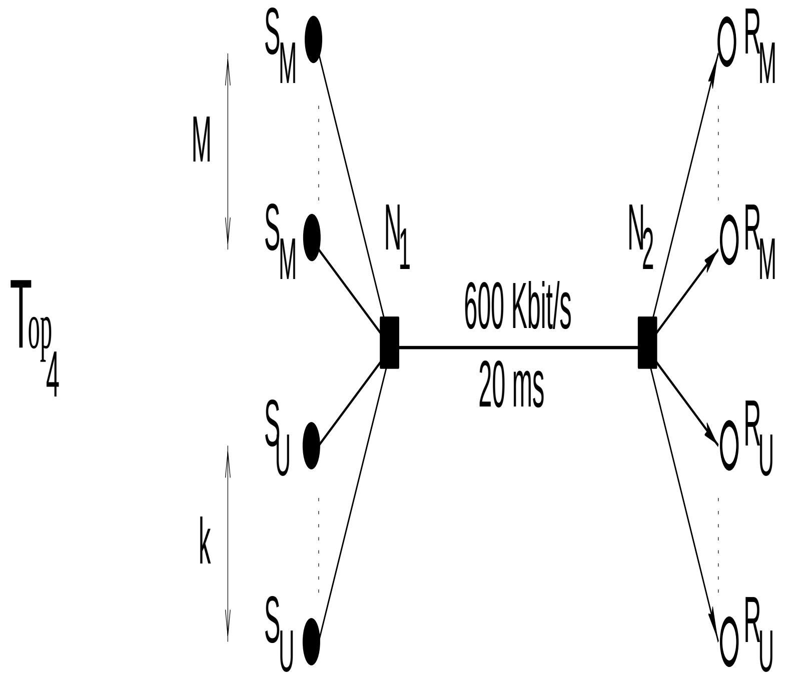

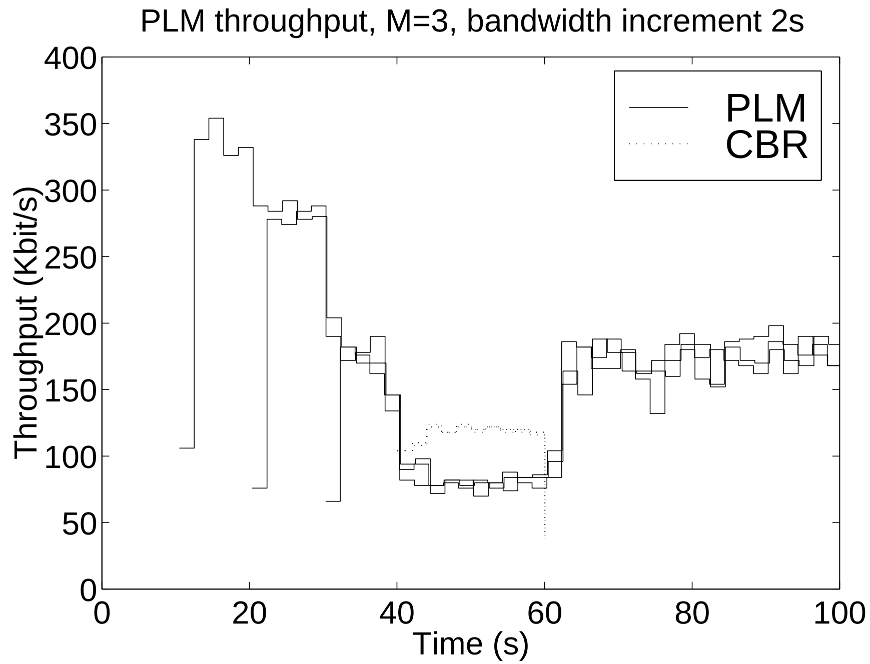

- Evaluation of the scaling of PLM with the number of sessions (mix of unicast and PLM sessions).
- 20Kbit/layer.
- PLM adapt to the available bandwidth in less than a RTT. **No loss induced** even in case of high congestion.

Multiple RLM and CBR Sessions

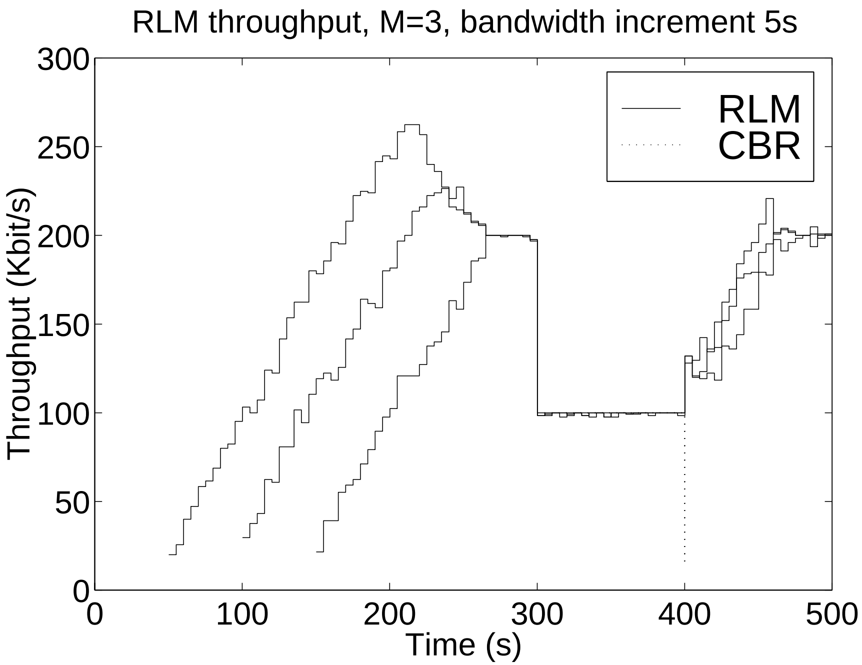

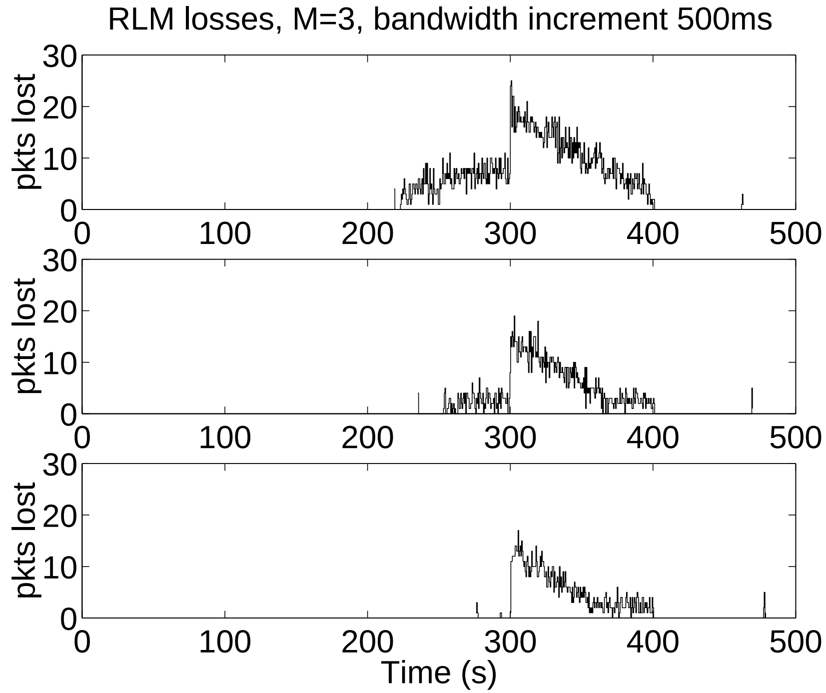

- Must run 500s. Low convergence time.
- High number of losses. Due to the conservative behavior of RLM (can not drop many layers in case of high congestion).

Simulations: Many PLM and TCP Flows

- We evaluate the behavior of PLM with an increasing number of PLM sessions and TCP flows. We considered many parameters (C, Layer granularity, Burst size).
- PLM+TCP: a realistic scenario.
- PLM performs incredibly well!

PLM and TCP: Scalability

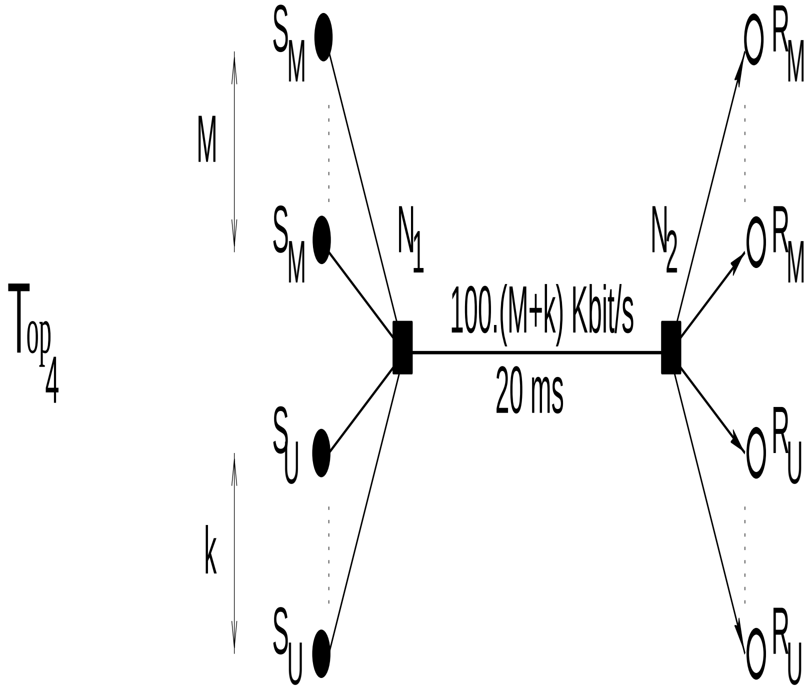

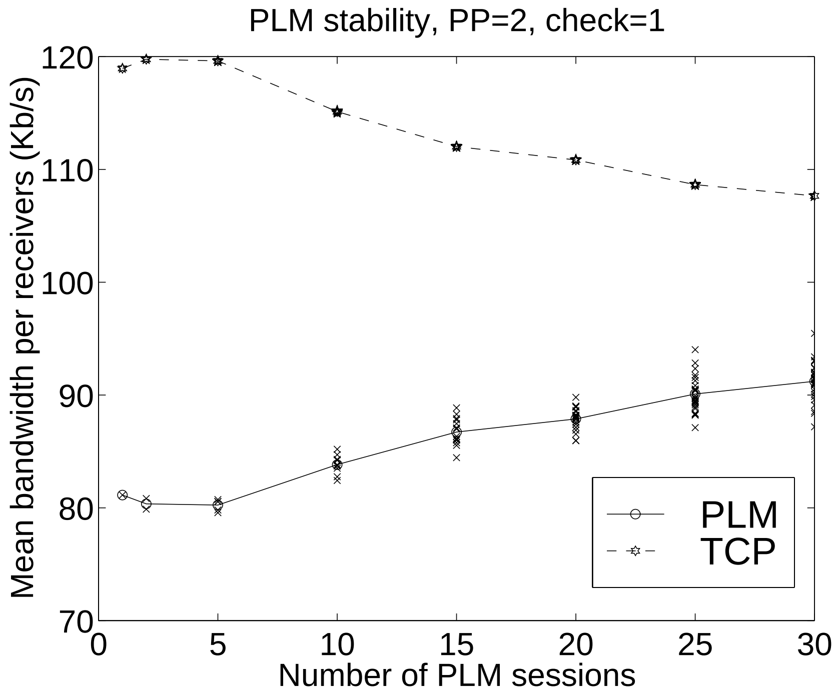

- Evaluation of the scaling properties of PLM with a large number of PLM and TCP flows.
- 20Kbit/layer.
- PLM gets all the available bandwidth (according to the layer granularity).

PLM and TCP: Scalability

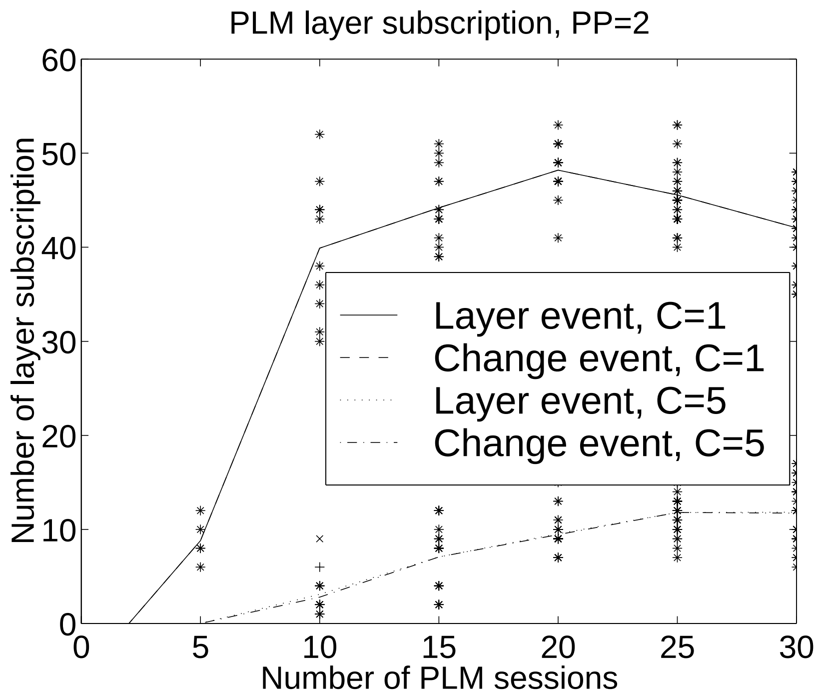

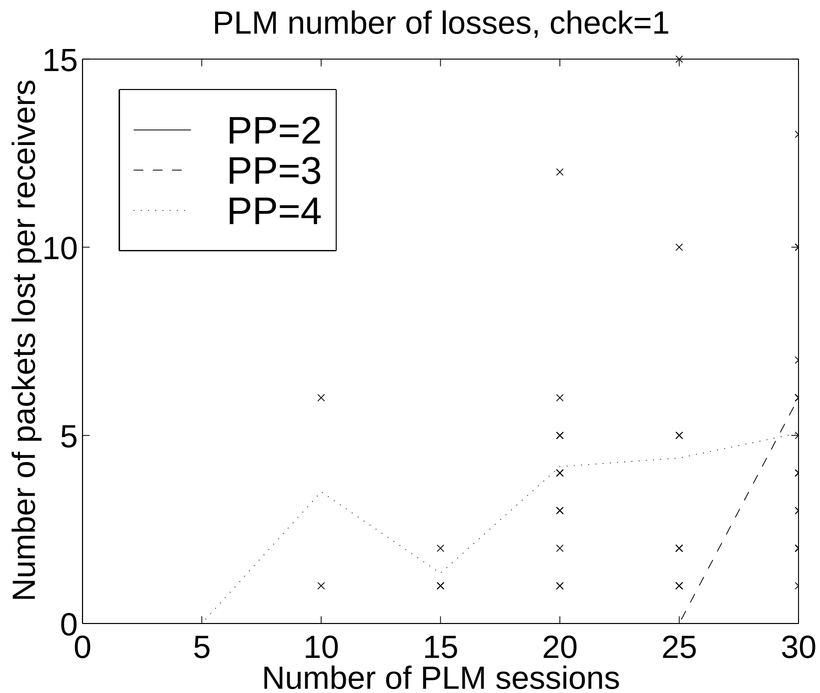

- Very low number of layer oscillations.
- **No loss induced** for PP=2, negligible number of losses for PP=3 and PP=4.

PLM vs. RLM: Parameters

- PLM parameters:
  - Check value C.
  - Burst size.
  - Layer granularity.
- RLM parameters:
  - Join-timer backoff constant.
  - Join-timer relaxation constant.
  - Detection-time estimator scaling term (2 parameters).
  - Detection-time estimator filter constants (2 parameters).
  - Loss threshold.
  - Maximum join-timer (frequency of the join experiment at equilibrium).
  - Minimum join-timer (start-up).
  - Etc.
- Parameters choice: still unresolved.

Final Conclusion

- We have defined a paradigm (the FS-paradigm) for the design of CC protocol. The FS-paradigm has appealing properties that convince us to devise a new CC protocol using this paradigm.
- We devise PLM, a new multicast CC protocol for audio/video dissemination:
  - PLM is a new CC protocol for multicast dissemination of audio/video content.
  - PLM outperforms all the previous CC protocols for audio/video.
  - PLM converges to the optimal rate in the order of one C and tracks this rate with no loss induced.
  - PLM is incontestably a practical validation of the FS-paradigm.